ROUTING ARCHITECTURE

FIELD OF THE INVENTION

This invention relates to the field of telecommunications, and more particularly to data communications.

BACKGROUND OF THE INVENTION

Data communication is a reflection of life in the 21st century. Applications, such as e-mail and the Internet, has increasingly become mainstream. Moreover, a move is afoot for migrating voice traffic from circuit-switched type networks to packet-switched type networks in support of Voice over IP ("VoIP") applications. Consequently, data traffic has continued to increase as acceptance and adoption of these applications continue to grow.

With the continued expansion of data applications, there is a growing consumer demand for accurate wired and wireless high-speed access. Systems supporting data communication typically employ a number of Application Processing ("AP") or communication nodes. These AP nodes may be driven by wired and wireless high-speed access, as well as VoIP applications.

AP nodes are interconnected through a transport or interconnect fabric fabric for the transmission of information therebetween. To support high-speed data communication, these interconnect fabrics may be cell or packet based to enable any one of a number of distinct high-speed data communication formats. Consequently, the routing or forwarding of cell or packet information has become an increasingly critical function.

Each AP node is typically realized by a circuit board. Within each AP node, cell or packet information may be routed or forwarded to any number of on-board processing devices by means of a dedicated switch. This switch effectively manages information traffic flow for the AP node's circuit board.

While the use of a dedicated switch is effective, there are notable limitations. Firstly, the dedicated switch consumes a non-trivial amount of power. Consequently, power consumption and heat dissipation issues may require attention. Moreover, the cost of each dedicated switch, and the space each consumes on the AP node's circuit board may also impact on the design and efficiency of the system.

Therefore, a need exists for an AP node architecture that avoids the limitations of the dedicated switch. Moreover, a routing architecture is desired that is supportive of improved power and cooling budgets, reduced board space consumed and overall cost.

SUMMARY OF THE INVENTION

The present invention provides a routing architecture for improving power and cooling budgets, as well as reducing board space consumed and overall cost. More particularly, the present invention provides a communication node architecture for routing cell and/or packet information between application processors. The present invention realizes the communication node architecture without the need for a dedicated switching device, such as an Ethernet switch, for example. The communication node architecture may be deployed in numerous applications, including, for example, a radio node controller, base station controller, and a traffic-processing controller.

In one embodiment, the communication node architecture of the present invention includes at least two network-processing devices for routing or forwarding cell and/or packet information, instead of the dedicated switching device known in the art. Each of the network processing devices may be interconnected by way of a shared bus structure, such as, for example, a Peripheral Component Interconnect ("PCI") bus. The shared bus structure may also couple the network processing devices with a general-purpose processing device, which

controls the duties performed by each of the network processing devices. At least one of the network processing devices may be coupled with a fabric for interconnecting one node with other nodes. The one or more network processing devices may be coupled to a fabric through a system interface. Additionally, cell and/or packet information may be received through a maintenance interface. It should be noted that the cell and/or packet information received by the maintenance interface (e.g., Operations and/or Maintenance type information) might be distinct over that received by the system interface (e.g., Bearer Transport Path Processing and/or Call Control type information). Consequently, the network processing device(s) may receive cell and/or packet information from the fabric through the system interface for routing or forwarding within the node.

In a further embodiment, the node architecture may also employ a multiplexer. The multiplexer may be used for coupling the network processing device(s) with the system interface and/or the maintenance interface. By this arrangement, the network processing device(s) may receive cell and/or packet information, multiplexed, from the maintenance interface and through the interconnect fabric by means of the system interface.

In still another embodiment, the one or more network processing device(s) may be coupled with an external system input/output through an interface device. The interface device may support one or more transport mechanisms. Consequently, the interface device may be designed to support, for example, Asynchronous Transfer Mode ("ATM"), Internet Protocol ("IP"), and/or Frame Relay ("FR").

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will be better understood from reading the following description of non-limiting embodiments, with reference to the attached drawings, wherein below:

FIG. 1 depicts a digital communication system;

FIG. 2 depicts an Applications Processing or communication node architecture; and

FIG. 3 depicts an embodiment of the present invention.

It should be emphasized that the drawings of the instant application are not to scale but are merely schematic representations, and thus are not intended to portray the specific dimensions of the invention, which may be determined by skilled artisans through examination of the disclosure herein.

DETAILED DESCRIPTION

Referring to **FIG. 1**, a high-level block diagram of a digital communication system **10** is illustrated. As depicted, system **10** is capable of supporting various services, including the communication of voice and/or data traffic. More particularly, system **10** may enable wired and/or wireless communication applications.

To further serve these purposes, digital communication system **10** includes a plurality of Application Processing ("AP") or communication nodes **20**. Each of the nodes **20** perform a function(s) required by system **10** and may be realized by a printed circuit board ("PCB"). The functions performed may include, for example, Bearer Transport Path Processing, Call Control, Operations and Maintenance. AP nodes **20** may be configured to execute the same function(s), and therefore may be realized using the same or similar design. In the alternative, each AP node **20** may serve a differing function(s), in part or in whole, depending on the requirements of system **10**, and therefore may be have relatively distinct configurations.

System **10** also includes a plurality of Interface Processor ("IP") nodes **40**. Each IP node **40** may interface with an external input/output port **45**. Moreover, each Interface Processor node **40** may also perform some processing on an incoming information stream to/from external input/output port **45**.

Digital communication system **10** may also include a transport or interconnect fabric **30** for enabling the transmission of information between AP nodes **20**. Using any one of a number of distinct communication formats, interconnect fabric **30**, in conjunction with system **10**, may be cell and/or packet based to support high-speed communication. As a result, the routing or forwarding of cell and/or packet information within system **10** is becoming an increasingly critical function.

Each AP node **20**, to this end, may be interconnected with one another through interconnect fabric **30**. Depending on the type of interconnect, a dedicated fabric card **50** may be required. Interconnect fabric **30** may be realized using an interconnect format type, depending on the number of applications supported by communication system **10**. The available interconnect format types may include, for example, a Time Division Multiplexed ("TDM") bus for use in circuit switched applications, Cell Based Interconnect for use in ATM applications, and/or Ethernet connectivity for use with packet switched applications. It should also be noted that digital communication system **10** may also support multiple interconnect format types, simultaneously, as well as a hybrid interconnect format therefrom.

Referring to **FIG. 2**, a block diagram of an exemplary architecture for an Application Processing ("AP") or communication node **100** is illustrated. AP node **100** performs one or more functions within the context of an application, such as, for example, digital communication system **10** of **FIG. 1**.

To support high-speed data communication, AP node **100** may employ multiple processors. As shown, AP node **100** includes a first and a second network processor, **110** and **120**, at least one of which receives cell and/or packet information. Moreover, AP node **100** also comprises a general-purpose processor **130**, which along with network processors, **110** and **120**, are coupled with each other, as well as an interconnect fabric (not shown) by means of a dedicated switch **140**. As shown, dedicated switch **140** may be realized by an

Ethernet switch, for example, particularly where the interconnect fabric is provided using an Ethernet based scheme.

Ethernet switch **140** performs several functions as part of AP node **100**. Firstly, Ethernet switch **140** provides an interconnection between multiple "ports" on AP node **100**, including redundant interfaces to an interconnect fabric, including a system interface **150**, as well as a maintenance interface **160**. Moreover, switch **140** provides an interconnection between multiple "ports" on AP node **100** and processors **110**, **120** and **130**.

Moreover, Ethernet switch **140** also functions as a routing device. Here, switch **140** routes and forwards Ethernet packets between ports. These ports may typically be based on an L2 (Ethernet) or an L3 (Internet protocol) set of routing instructions. It should be noted that general-purpose processor **130** may act here as a traffic control mechanism for switch **140.**

Ethernet switch **140** may also provide a fail-over feature for AP node **100**. Here, switch **140** may assist in handling the fail-over of redundant interfaces to an interconnect fabric, such as system interface **150**. Moreover, switch **140** may switch from an active port to a standby port in the event of the detection of a failure on the active port.

It should be noted that Ethernet switch **140** may perform other functions, as called for by AP node **100**. The functions may be of particular relevance given the application of the AP node **100** in the context of its application in a digital communication system. These functions may include buffering, support for Class of Service, and flow control, for example.

While the use of Ethernet switch **140** within AP node **100** serves several beneficial purposes, notable limitations remain. Firstly, the dedicated switch consumes a significant amount of power. Consequently, power consumption and heat dissipation issues may require attention. Moreover, the cost of each dedicated switch, and the space each consumes on the AP node's circuit board may also impact on the design, capacity and efficiency of the system.

Referring to **FIG. 3**, an embodiment of the present invention is illustrated. More particularly, a routing architecture **200** is depicted for addressing the limitations associated with using dedicated (e.g., Ethernet) switch **140** within AP node **100** of **FIG. 2**. Routing architecture **200** obviates the need for a dedicated switch device in the architecture, in favor of a more distributed approach.

It should be noted that the flexible nature of routing architecture **200** might enable an AP node to support the native transport of multiple cell or packet protocols simultaneously. This added flexibility may allow an AP node(s) to address additional applications previously not possible in the known art. In this regard, routing architecture **200** may simultaneously route and/or forward a cell(s) and/or a packet(s) (e.g., Ethernet, IP, ATM), in parallel, for example, within architecture **200**.

Routing architecture **200** provides a superset of the capabilities without requiring a dedicated switch element for performing cell and/or packet routing and/or forwarding. To this end, routing architecture **200** receives cell and/or packet information through an interconnect fabric **210**. Interconnect fabric **210** couples the AP node, as reflected in routing architecture **200**, with another AP node (not shown). It should be noted that in the present disclosure, while reference is made to routing architecture **200** receiving cell and/or packet information from interconnect fabric **210**, cell and/or packet information may also be transmitted to interconnect fabric **210** after being processed by the components forming routing architecture **200**, disclosed hereinbelow. Consequently, for simplicity, reference to term receiving herein may include transmitting.

Received cell and/or packet information may be fed into or out of routing architecture **200** by means of a system interface **220**. Cell and/or packet information may also be received by a maintenance interface **230**. In one embodiment, the cell and/or packet information received by maintenance interface **230** may correspond with Operations and/or Maintenance type

information, for example. In contrast, the cell and/or packet information received by system interface **220** may correspond with Bearer Transport Path Processing and/or Call Control type information, for example.

To process the aforementioned cell and/or packet information, routing architecture **200** includes a plurality of network processing ("NP") devices, **240**, **250** and **260**. More particularly, one or more NP devices, **240**, **250** and/or **260**, may be designated for receiving cell and/or packet information from interconnect fabric **210** by means system interface **220**. Thusly, system interface **220** may couple at least one NP device, **240**, **250** and/or **260**, with the fabric **210** to facilitate communication between distinct AP nodes.

To support the functionality assumed by NP devices, **240**, **250** and **260**, routing architecture **200** may also include a shared bus structure **270**. Shared bus structure **270** provides a means for coupling each of NP devices, **240**, **250** and **260**, with one another on the same AP node corresponding with routing architecture **200**. In one embodiment, shared bus structure **270** may comprise a Peripheral Component Interconnect ("PCI") bus.

Routing architecture **200** also may include a general-purpose processor **280**. General-purpose processor **280** may serve a multitude of functions, including controlling each of NP devices, **240**, **250** and **260**. Moreover, general-purpose processor **280** may also perform maintenance on the AP node, as realized by routing architecture **200**. In support of these functions, general-purpose processor **280** may also be coupled with shared bus structure **270**.

By the above configuration, NP devices, **240**, **250** and/or **260**, may also perform additional functions. One or more NP devices, **240**, **250** and/or **260**, for example, may determine the destination of the received cell and/or packet information within routing architecture **200**. In one embodiment, the destination of the cell and/or packet information may be determined in response to one or more stored routing rules and/or particular characteristics of the cell and/or packet information (e.g., packet type, L2, L3, destination address,

source address and other packet information). Thereafter, at least one NP device, **240**, **250** and/or **260**, may forward or route the cell and/or packet information to the determined destination.

It should be noted that one or more NP devices, **240**, **250** and/or **260**, may support peer-to-peer routing. Peer-to-peer routing here may mean routing between one NP device, **240**, **250** or **260**, and one or more other NP devices, **240**, **250** and/or **260**. Similarly, peer-to-peer routing may also include routing between general-purpose processor **280** and one or more NP devices, **240**, **250** and/or **260**.

Routing architecture **200** may also support a direct delivery feature. Here, a cell(s) and/or packet may be delivered directly from general-purpose processor **280** or one NP device, **240**, **250** or **260**, into the memory of one or more other processing devices (e.g., another NP device(s), **240**, **250** and/or **260**, and/or general-purpose processor **280**) via the shared bus structure **270**, for example. By this arrangement, the delivered cell(s) or packet(s) may arrive without interrupting (or waking these one or more other processing devices), which may be processing other information (or sleep mode operation) at the time. Consequently, when one or more of these other processing devices are ready (or awoken), the specific cell(s) or packet(s) is waiting to expedite its subsequent internal processing. In the alternative, the specific cell(s) or packet(s) may arrive directly into the memory of one or more other processing devices, thereby initiating an interrupt (or a wake up) routine therein.

It should also be noted that the routing and/or forwarding of the cell and/or packet information between NP devices, **240**, **250** and/or **260**, and/or general-purpose processor **280** relies on aspects of programmability to exploit the flexibility structure of architecture **200**. In this regard, the stored routing rules may vary from simple to complex. However, these the routing rules may be embodied in software, and thusly updated and/or upgraded to provide even greater flexibility.

Routing architecture **200** may also allow for deep packet inspection at the NP device level. Here, deep packet inspection may afford architecture **200** the ability to give routing software access to some or all of the fields in the cell and/or packet. Routing and/or forwarding may then be performed using a single field or a plurality of fields obtained from a plurality of protocol layers in the packet, for example. For example, Ethernet packets with the same layer 3 destination IP address may be delivered to different NP devices, **240**, **250** and/or **260**, based on layer 4 and higher parameters (e.g., UDP port number).

In one embodiment, routing architecture **200** also comprises a multiplexer **290**. Multiplexer **290** couples system interface **220** and maintenance interface **230** with one or more of NP devices, **240**, **250** and/or **260**. By this arrangement, multiplexer **290** creates a multiplexed stream from cell and/or packet information received of system interface **220** and of maintenance interface **230** to enable at least one NP device, **240**, **250** and/or **260**, to perform various process functions detailed herein.

Routing architecture **200** may also comprise at least one external system input/output interface **300**. External system input/output interface **300** may be coupled with one or more NP devices, **240**, **250** and/or **260**. As an external system input/output interface **300**, support for one or more transport mechanism types may be required. Consequently, external system input/output interface **300** may support at least one of Asynchronous Transfer Mode, Internet Protocol, and Frame Relay, for example.

EXEMPLARY EMBODIMENTS

In an AP node based on architecture of the present invention, the routing and/or forwarding of cells or packets should be performed by one of the NP devices. In a typical application of the AP node, this routing and/or forwarding functionality may require only a subset of the resources of the NP device(s). As

a result, an unidentified amount of resources may remain available in the NP device(s) for other processing.

Referring to the embodiment of **FIG. 3**, at least NP device **240**, for example, may receive packets or cells coming from system interface **220** or maintenance interface **230**, possibly via through optional multiplexer **290**. At least one NP device **240** should determine a destination, such as another NP device **250** and/or NP device **260** and/or general-purpose process **280**, for a given cell and/or packet and move the cell and/or packet to that designated device(s) via the Shared Bus or some other convenient interconnection means. If the destination for a given cell and/or packet is NP device itself, the packet should be forwarded locally. The process in which the cell and/or packet information is moved should be based on a set of locally stored routing rules, and possibly other characteristics of the cell or packet, such as the source port.

In the reverse direction, at least NP device **240**, for example, may collect cells and/or packets from the remaining processors via the Shared Bus. Consequently, at least NP device **240**, for example, may then forward cells and/or packets to the appropriate port based on locally stored routing rules. Here, at least NP device **240** may also be capable of supporting peer-to-peer routing between processors.

It can be seen that the important functions previously performed by a dedicated switch, such as Ethernet switch **140** of **FIG. 2**, may be distributed between the routing software in at least one network processor, such as NP device **240** of **FIG. 3**, for example, along with a shared bus structure, and possibly an optional multiplexer. In this regard, shared bus structure and the optional multiplexer may handle the interconnection between the various "ports" on the board, including the redundant system interfaces, the local maintenance interface, and the multiple processors. Moreover, the forwarding and/or routing of cell and/or packets may be performed by a subset of the processing resources in one of the NP devices. Similarly, handling the fail-over

of the redundant system interfaces may be performed by a combination of one of the NP devices and the optional multiplexer. Finally, functionality such as buffering, support for Class of Service, and flow control may also be performed by a subset of the processing resources in an NP device.

It should be noted that without the dedicated switch, such as Ethernet switch **140** of **FIG. 2**, the routing architecture of the present invention may exhibit improved performance in terms of power budget, heat dissipation, board space, and cost. Given these enhancements, it is possible to design an AP node with more processor elements, further boosting system performance. In addition to the performance enhancements outlined above, this architecture exhibits enhanced flexibility, which will allow Application Processors designed with this architecture to address new applications.

By eliminating the need for the dedicated switch in the architecture, it is possible for an AP node to perform additional functions. For example, as a result of the present invention, an AP node may be able to simultaneously support multiple cell or packet transport protocols, as well as their transfer therein. This may be attributed to the cell or packet routing and/or forwarding mechanism implemented in a programmable element of one of the processors on the board.

The flexibility of the routing architecture of the present invention has several advantages. First, an AP node utilizing this architecture may support multiple cell or packet transport mechanisms, where the cells and packets are transported in their native format. This may provide performance enhancements over a system using encapsulation to support multiple formats. Secondly, an AP node in accordance with the present invention may be capable of supporting additional applications. It can easily be seen that an AP node may be configured to support applications requiring interfaces to external system input/output, such as an Interface Processor, for example.

Furthermore, the routing architecture of the present invention may also support pre-stripping of packet header information. More particularly, a packet(s) may be routed and/or forwarded amongst processing devices forming the AP node without the need to utilize header information. This is in contrast with dedicated switch of other architectures used previously (e.g., Ethernet switch **140** of **FIG. 2**).

It should be noted that NP devices, **240**, **250** and/or **260**, shared bus structure **270**, and general-purpose processor **280**, may be each be configured in furtherance of the flexibility of routing architecture **200**. As stated hereinabove, routing architecture **200** may simultaneously and/or concurrently route and/or forward a cell(s) and/or a packet(s). Consequently, architecture **200** cell(s) and/or packet(s) data may be moved between NP devices, **240**, **250** and/or **260**, shared bus structure **270**, and/or general-purpose processor **280**.

In another embodiment of the present invention, if the routing and/or forwarding mechanisms detailed herein were operating on as Ethernet packets, an ATM Interface block may be coupled with an external system input/output for transporting ATM cells over some appropriate physical interface. Here, the ATM Interface block may also connected to one of the network processors, where the routing and/or forwarding mechanisms may be implemented. It may be advantageous in this example for the connection means to the network processor to be realized by a Utopia bus, for example.

Operation of the routing architecture of the present invention would be as follows. Initially, one of the network processors may be programmed to support the additional transport means (e.g., Utopia), and be used for an additional set of locally stored routing rules to determine the appropriate destination processor for an incoming cell. Here, the identified network processor may then forward the cell over the shared bus structure or the like. If the cell is destined for the one network processor itself, the cell may be forwarded locally.

In the reverse direction, the one network processor may collect cells from the other processors via the Shared Bus. After examining the locally stored routing rules, the one network processor may determine that the cells are destined for the ATM Interface block. Thereafter, the one network processor may enable the cells to be forwarded over the Utopia bus, for example.

While the particular invention has been described with reference to illustrative embodiments, this description is not meant to be construed in a limiting sense. It is understood that although the present invention has been described, various modifications of the illustrative embodiments, as well as additional embodiments of the invention, will be apparent to one of ordinary skill in the art upon reference to this description without departing from the spirit of the invention, as recited in the claims appended hereto. Consequently, processing circuitry required to implement and use the described system may be implemented in application specific integrated circuits, software-driven processing circuitry, firmware, programmable logic devices, hardware, discrete components or arrangements of the above components as would be understood by one of ordinary skill in the art with the benefit of this disclosure. Those skilled in the art will readily recognize that these and various other modifications, arrangements and methods can be made to the present invention without strictly following the exemplary applications illustrated and described herein and without departing from the spirit and scope of the present invention. It is therefore contemplated that the appended claims will cover any such modifications or embodiments as fall within the true scope of the invention.